

Autonomous Navigation Performance During The Hartley 2 Comet Flyby

Matthew J Abrahamson¹, Brian A. Kennedy², Shyam Bhaskaran³

Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA, 91109-8099

On November 4, 2010, the EPOXI spacecraft performed a 700-km flyby of the comet Hartley 2 as follow-on to the successful 2005 Deep Impact prime mission. EPOXI, an extended mission for the Deep Impact Flyby spacecraft, returned a wealth of visual and infrared data from Hartley 2, marking the fifth time that high-resolution images of a cometary nucleus have been captured by a spacecraft. The highest resolution science return, captured at closest approach to the comet nucleus, was enabled by use of an onboard autonomous navigation system called AutoNav. AutoNav estimates the comet-relative spacecraft trajectory using optical measurements from the Medium Resolution Imager (MRI) and provides this relative position information to the Attitude Determination and Control System (ADCS) for maintaining instrument pointing on the comet. For the EPOXI mission, AutoNav was tasked to enable continuous tracking of a smaller, more active Hartley 2, as compared to Tempel 1, through the full encounter while traveling at a higher velocity. To meet the mission goal of capturing the comet in all MRI science images, position knowledge accuracies of ± 3.5 km (3- σ) cross track and ± 0.3 seconds (3- σ) time of flight were required. A flight-code-in-the-loop Monte Carlo simulation assessed AutoNav's statistical performance under the Hartley 2 flyby dynamics and determined optimal configuration. The AutoNav performance at Hartley 2 was successful, capturing the comet in all of the MRI images. The maximum residual between observed and predicted comet locations was 20 MRI pixels, primarily influenced by the center of brightness offset from the center of mass in the observations and attitude knowledge errors. This paper discusses the Monte Carlo-based analysis that led to the final AutoNav configuration and a comparison of the predicted performance with the flyby performance.

I. Introduction

On July 4, 2005, the Deep Impact Flyby spacecraft flew by the comet Tempel 1 and observed the impact of the Impactor spacecraft with Tempel 1. Following this successful prime mission, the Flyby spacecraft was repurposed for the EPOXI extended mission. The highest priority of the EPOXI science observations was a low-altitude flyby of the comet Hartley 2 to capture high-resolution images of the nucleus. A major challenge of the mission was determining how best to meet new mission requirements with an existing spacecraft while applying minimal engineering modifications.

On November 4, 2010, the EPOXI spacecraft flew by Hartley 2 at an altitude of 700 km and a relative velocity of 12.3 km/s. The flyby of Hartley 2 was the fifth time that a comet nucleus has been imaged by a spacecraft. Of the five comets visited, Hartley 2 is by far the smallest comet nucleus imaged with the most dynamic activity per surface area. Prior comet nucleus imaging was performed by the Giotto spacecraft at Halley, the Deep Space 1 spacecraft at Borrelly², the Stardust spacecraft at Wild 2³, and the Deep Impact spacecraft at Tempel 1^{4,5}. Of these missions, all except the Halley flyby contained an autonomous navigation software package onboard, called AutoNav, to aid science collection. AutoNav's purpose is to perform orbit determination on the flight computer

¹ Navigation and Mission Design Section, Jet Propulsion Laboratory, California Institute of Technology, MS 301-121, 4800 Oak Grove Drive, Pasadena, CA, 91101-8099; Matthew.Abrahamson@jpl.nasa.gov, 818-354-2880.

² Navigation and Mission Design Section, Jet Propulsion Laboratory, California Institute of Technology, MS 264-820, 4800 Oak Grove Drive, Pasadena, CA, 91101-8099; Brian.M.Kennedy@jpl.nasa.gov, 818-354-6327.

³ Navigation and Mission Design Section, Jet Propulsion Laboratory, California Institute of Technology, MS 264-820, 4800 Oak Grove Drive, Pasadena, CA, 91101-8099; Shyamkumar.Bhaskaran@jpl.nasa.gov, 818-354-3152.

using optical measurements from the spacecraft's main camera. For the purposes of comet flybys, it provides a capability to continually update the onboard estimate of the spacecraft-comet relative trajectory as stronger observations become available at closer range. Without the use of AutoNav, science observations would have to be captured as a mosaic by scanning the camera across the sky to capture the uncertainty in the knowledge of the comet position.

The high velocity, low altitude profile of the Hartley 2 flyby required highly accurate knowledge of the comet-relative spacecraft trajectory in order to maintain camera pointing on the comet throughout the encounter. The primary camera on the spacecraft, the Medium Resolution Imager (MRI), has a 10-milliradian field of view, requiring comet-relative trajectory knowledge of ± 3.5 km ($3\text{-}\sigma$) to maintain visual lock at the 700-km flyby altitude. Given the flyby velocity of 12.3 km/s, this translates to a time of flight knowledge requirement of ± 0.3 seconds ($3\text{-}\sigma$). For the EPOXI flyby, the trajectory knowledge requirements stated above could not be met by ground-based navigation. The final ground-based trajectory uplinked to the spacecraft at 24 hours prior to the flyby was delivered with a formal sigma of ± 3.1 seconds ($3\text{-}\sigma$) on the time of flight, exceeding the requirements by a factor of 10. Trajectory reconstruction of the flyby indicated an actual time of flight error of -1.05 seconds in the ground-based trajectory, still exceeding the requirements by more than a factor of three. To meet the requirements, AutoNav was enabled to begin capturing optical measurements at 50 minutes prior to closest approach, and to begin updating the onboard ephemeris at 42 minutes prior to closest approach. By the ephemeris update at six minutes prior to closest approach, AutoNav met all of the above trajectory knowledge requirements, ensuring comet capture in the MRI throughout the encounter. The updated ephemeris solutions produced by AutoNav were used by ADCS to maintain camera pointing on the comet.

In preparation for the flyby, the AutoNav configuration was tuned for robust system performance when presented with trajectory, measurement, and attitude knowledge errors. A series of Monte Carlo simulations assessed AutoNav's statistical performance when subjected to these disturbances and provided a testbed for validating various operational configurations. This analysis influenced key operational decisions including critical AutoNav parameters (such as start time, measurement arc length, and filter weighting strategy) and spacecraft parameters (such as attitude mode transition times and time of return to the ground-based orbit solution on departure).

The AutoNav performance at Hartley 2 was successful, capturing the comet in all of the MRI images. The maximum residual between predicted and observed camera frame locations was 20 pixels, primarily influenced by the comet's center of brightness offset from its center of mass in the observations and the star tracker bias in the attitude estimate. This paper discusses the Monte Carlo simulation analyzing AutoNav performance and its influence in determining the final AutoNav configuration. Comparisons between simulation and flight performance will also be discussed.

II. Prior Autonav Usage

The original AutoNav system was developed at JPL in the late 1990s to fly on the technology demonstration mission, Deep Space 1¹. On that mission, AutoNav exercised its capacity to perform autonomous optical orbit determination in cruise, maintain the orbit using ion propulsion and conventional thrusting, and perform autonomous optical orbit determination during a small body flyby. On September 23, 2001, the Deep Space 1 spacecraft flew by the short period comet Borrelly at a distance of approximately 2200 km². The AutoNav system was enabled 30 minutes prior to closest approach and maintained lock on the comet nucleus until 2 minutes before closest approach, capturing the comet nucleus at a resolution of 40 meters per pixel.

On January 2, 2004, the Stardust spacecraft flew by the comet Wild 2 at a distance of 236 km to capture a sample of comet dust for return to Earth³. On this mission, AutoNav was used to support flyby imaging of Wild-2. Stardust collected the images using a gimbaled mirror system, which directed the camera field of view toward the comet position determined by AutoNav without slewing the entire spacecraft platform. The AutoNav version flown on Stardust contained the ability to estimate three components of position and three components of attitude, but did not estimate velocity³. AutoNav was enabled 20 minutes prior to closest approach and returned images of Wild 2 at a resolution of 14.5 meters per pixel.

The version of AutoNav onboard the Deep Impact spacecraft was designed and tested for the Deep Impact mission at comet Tempel 1 on July 4, 2005⁴. In that mission, the Deep Impact Flyby spacecraft separated from the Impactor spacecraft 24 hours prior to the Tempel 1 flyby. Each spacecraft executed an independent AutoNav system with image processing, state estimation, and, for the Impactor spacecraft, maneuver calculation capabilities. The AutoNav version flown on Deep Impact provided the capability to estimate three components of position, three components of velocity, and two components of attitude rate error, but no capability to estimate attitude bias error⁴.

While the AutoNav system onboard the Impactor determined a target impact location and commanded maneuvers to reach that location, the Flyby spacecraft independently estimated the impact location as well as its own trajectory. The objective of the Flyby spacecraft was to observe the impact and resulting plume while protecting sensitive science instruments from the comet environment. The Deep Impact mission was a phenomenal success, returning images of the comet nucleus and impact at a resolution of 7 meters per pixel⁵.

Part of the challenge of designing the EPOXI mission was repurposing the existing Deep Impact spacecraft for the new requirements of the Hartley 2 flyby. This included applying the existing flight version of AutoNav to a new navigation problem involving new dynamics and new observation challenges.

III. Spacecraft Flight System

The Deep Impact Flyby Spacecraft system was designed and built by Ball Aerospace Technologies Corporation for remote comet observation. It features high-resolution visible and infrared instruments coupled with a precise attitude control system and autonomous navigation capability to return comet nucleus observations. It communicates to the ground using either a High-Gain Antenna or a Low-Gain Antenna and contains two RAD750 computers for processing. The spacecraft bus is mounted to the back of a large solar array, which provides system power and doubles as a dust shield.

The science instruments resident on the spacecraft are the Medium Resolution Imager (MRI) and the High Resolution Imager (HRI). The MRI captures visible images on a 1024x1024 CCD with a 10 mrad field of view. It features a 12 cm aperture, a focal length of 2100 mm, and a pixel scale of 10 μ rad per pixel. The HRI captures visible images using the same CCD electronics as the MRI with a 30 cm aperture, 10500 mm focal length, and 2 μ rad per pixel scale. Although the HRI features a resolution five times greater than the MRI, a focusing problem discovered shortly after launch limits its performance to a level similar to the MRI. Also contained on the HRI instrument is an infrared spectrometer (IR) that scans a wavelength spectrum from 1.05 to 4.8 microns within a 0.29-degree (5.1 mrad) slit^{4,5}.

The MRI and HRI are fixed to the spacecraft bus platform and rely on attitude maneuvers of the entire spacecraft bus for accurate pointing. The ADCS system achieves three-axis stabilized control using momentum wheels, with four hydrazine-fueled RCS thrusters available for desats and trajectory correction maneuvers (TCMs). The ADCS system also includes two star trackers and an Inertial Reference Unit to estimate the attitude of the spacecraft^{4,5}.




Figure 1. Deep Impact Spacecraft

IV. Mission Objectives

The EPOXI mission exercised the existing spacecraft in a new set of operating regimes. Table 1 compares several attributes of the Tempel 1 and Hartley 2 flybys. The major objective of EPOXI was to capture images of Hartley 2 continuously from approach through departure, including at closest approach. In contrast, the previous flyby at Tempel 1 ceased tracking of the comet 50 seconds prior to closest approach and did not resume tracking on departure until 40 minutes after closest approach^{4,5}. The gap in science data allowed for an attitude profile that protected the instruments from dust and downlinked impact images to the ground, reflecting the higher priority of impact observations at Tempel 1. The EPOXI mission prioritized high-resolution comet nucleus observations, returning close range images from the frontside, underside, and backside of Hartley 2 by engaging AutoNav throughout the entire encounter. To track the comet through closest approach, the spacecraft was required to slew 180 degrees, with a maximum expected slew rate of 1 degree per second. These new conditions greatly tightened requirements on the ADCS and AutoNav subsystems.

In addition to the continuous tracking objective, EPOXI was also tasked to track a new and relatively unknown target. Hartley 2 represents a much smaller and more dynamically active comet as compared to Tempel 1. In fact,

Table 1. Comparison of Deep Impact and EPOXI Missions

	Deep Impact	EPOXI
Primary Objective	Observe impact	Observe nucleus
Observation Gap	E-50 sec to E+40 min	No Gap
Comet Size	7.6 km x 4.9 km	2.2 km x 0.5 km
Comet Period	40.7 hours	18.1 hours
Relative Velocity	10.2 km/s	12.3 km/s
Solar Phase Angle at Closest Approach	64°	77°

the semi-major axis of the Hartley 2 nucleus was more than a factor of two smaller than any previously imaged comet. Hartley 2 has also been described as a highly dynamic comet, with large plumes of material regularly ejected from the nucleus in the form of jets. Lastly, the dynamics of the flyby resulted in a 12.3 km/s relative velocity during Hartley 2 imaging, compared to a 10.2 km/s relative velocity at the Tempel 1 flyby^{4,5}.

The flyby trajectory was designed to target an altitude of 700 km and a solar phase angle of 77° at closest approach. A convenient method of specifying the closest approach conditions is via the B-plane. The B-plane is centered on the target body and perpendicular to the incoming asymptote of the trajectory. The intersection of the approach asymptote and the B-plane is the spacecraft position at the closest approach time. The components of the B-plane are referred to as $B \cdot T$ (horizontal projection parallel to a reference plane), $B \cdot R$ (vertical projection perpendicular to a reference plane). The reference plane used to build the EPOXI B-plane is the Ecliptic plane, meaning the $B \cdot T$ coordinate is parallel to the Ecliptic and contains the majority of the center of brightness offset signal. The third component of the B-plane frame is the linearized time-of-flight (LTOF), representing the time at which the trajectory asymptote crosses the B-plane. This system is especially useful for small body flybys since the target body is massless for all practical purposes and thus the B-plane coordinates represent a close approximation to the actual spacecraft position at the time of closest approach.




Figure 2. B-plane Coordinates

Figure 3 shows the closest approach target conditions for the EPOXI flyby. The targeted coordinates in the B-plane were $B \cdot R = 695.8$ km, $B \cdot T = -77.4$ km. In terms of the reference plane, EPOXI flew 695.8 km below Hartley 2 at closest approach and 77.4 km to the left. The relative velocity vector is into the page and the relative position vector points upward toward the origin of the plot. During the flyby, the camera attitude is configured to point along the spacecraft-to-comet position vector. In order to continuously track the comet through closest approach, the camera slews nearly 180° through the trajectory plane, or orbit plane, about the slew axis pointing to the right. The camera slews from looking forward along the asymptote on approach, to looking upward at closest approach, to looking backwards along the negative asymptote on departure. To simplify the projection of the dynamics in the camera frame, the spacecraft attitude aligns the camera axis with the trajectory plane throughout the slew.

Figure 4 depicts the projection of these axes in an MRI image. The spacecraft-to-comet vector points into the page, the slew axis points down, and the sun direction is down and slightly to the left. Since the axis of the nominal slew is vertical and the camera turns perpendicular to this axis, the nominal motion is in the horizontal direction.




Figure 3. Hartley 2 Flyby B-Plane

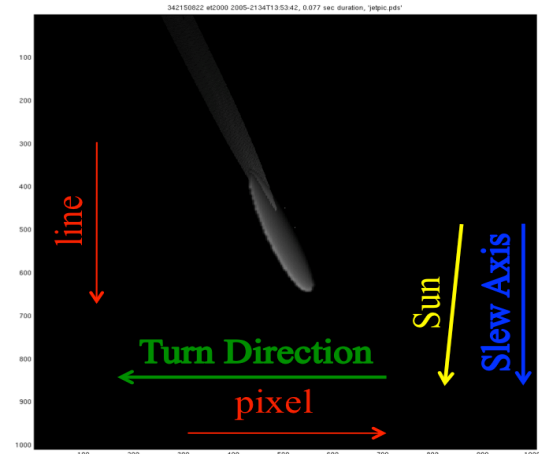


Figure 4. Geometry Projections In MRI Image

Locations in the image are measurement in pixel and line coordinates. The line coordinate increases from 0 to 1024 from top to bottom, and the pixel coordinate increases from 0 to 1024 from left to right. Given the projection of the

axes, the center of brightness offset will be mainly observable in the line axis, while the slew time of flight errors will be observable in the pixel axis.

V. AutoNav Overview

AutoNav is a software package developed for the purpose of performing autonomous orbit determination using optical measurements of a nearby target body. It provides a capability to update onboard knowledge of the spacecraft trajectory relative to a target body in near-real time. Its intended use is for close proximity navigation around small bodies, where ground-in-the-loop navigation is not practical due to long transmission and ground processing times. AutoNav is particularly useful for hyperbolic flybys, such as the Hartley 2 flyby, since the parallax necessary to resolve along-track trajectory error can only be observed within a few minutes of closest approach.

AutoNav's capabilities include image processing, state estimation, and maneuver computation, but only the first two capabilities were used for EPOXI. Image processing refers to the process of identifying regions of signal in a camera picture, determining the most likely comet signal, and producing an observed measurement of the comet relative to the spacecraft. AutoNav initiates image processing when a visual exposure is received from either the HRI or MRI instrument. The determination of the comet signal must be robust to disturbances by cosmic rays, hot pixels, image noise, bright outbursts from the comet, and other spurious signal in the image. Image processing provides two modes for determination of the comet signal in each image. The first mode, called Blobber, searches for all signal in the picture above a brightness threshold, segments it into contiguous areas of brightness, and returns the area and maximum brightness of each blob to the AutoNav executive. The centroid of the largest size blob is used for orbit determination. Blobber mode is preferred during the first sets of AutoNav imaging when ephemeris predictions, comet size, and comet albedo are uncertain and the comet is not bifurcated. The second mode, called Centroid, applies a moment algorithm to a defined subframe of the picture to calculate the centroid location of all pixels above a brightness threshold. The subframe is sized based on the range to the comet. This mode is used once the comet signal is well determined and automatically excludes any disturbances outside of the image subframe.

The observations determined by image processing are passed to a batch-sequential least-squares navigation filter for state estimation. The process of estimating and delivering a new trajectory model to the spacecraft is called an orbit determination (OD) update. The data arc of observations sent to the filter is determined by an OD arc length parameter configured for AutoNav. This parameter determines how many minutes of recent observation data is included in the data arc sent to the filter. As new measurements are incorporated into the data arc, the oldest measurements are removed. The filter determines the residuals between the predicted measurements and the observations from image processing for determining a least squares fit of the spacecraft trajectory relative to the target body. The state estimate is then passed to the ADCS for orientation of the science instruments. For the EPOXI flyby, Image Processing was triggered every 15 seconds and a new OD update was provided every minute, similar to the cadence used at Tempel 1^{4,5}.

Since the Deep Impact spacecraft was already flying with AutoNav onboard, a main objective of EPOXI was to reconfigure the AutoNav system for the new mission without modifications to the flight code. AutoNav configurations were analyzed using a "flight-code-in-the-loop" Monte Carlo simulation to predict AutoNav's statistical performance when subjected to various disturbances. The AutoNav team used this Monte Carlo simulation to assess the system performance given the new trajectory dynamics, new comet model, and new mission requirements.

VI. Attitude Determination and Control

The Attitude Determination and Control System is coupled with the AutoNav system to allow spacecraft tracking of remote objects in the MRI or HRI. The functions of ADCS are (1) to estimate the current pointing of the spacecraft and (2) to control it to point at the relative comet position, while it is AutoNav's function to determine the relative comet position. ADCS estimates the spacecraft attitude using a combination of gyroscope and star tracker data and controls the attitude using a set of momentum wheels. The ADCS control algorithm uses the AutoNav OD update to construct a new target attitude, angular rate, and angular acceleration for the spacecraft platform.

The ADCS attitude estimate is a critical input to AutoNav, since it provides the instrument pointing at the time of each image exposure. Errors in the attitude estimate, or nominal attitude relative to the true attitude, are defined as attitude knowledge errors. The presence of attitude knowledge errors can cause AutoNav to falsely interpret the optical measurements taken by the camera. In particular, attitude errors can be incorrectly estimated as position or velocity errors in the AutoNav filter. The flight version of AutoNav onboard the EPOXI spacecraft does not include the capability to estimate attitude bias errors, so a major undertaking of the AutoNav team was effective management of these errors within the constraints of the software.




Figure 5. AutoNav-ADCS System Interaction

ADCS attitude determination has three modes of operation for estimating the spacecraft attitude: (1) Nominal mode, (2) Override1 Mode, and (3) Override2 Mode. These modes determine which measurements are taken and how quickly they are incorporated into the attitude estimate. In Nominal Mode, both star tracker and gyroscope measurements are taken, with the star tracker data merged into the attitude solution with a 8000 second time constant. Override1 Mode only incorporates gyroscope data, while Override2 Mode incorporates both star tracker and gyroscope data with a 100 second time constant. For the EPOXI flyby, Override1 Mode is used to ignore the star trackers around closest approach when slew rates are too fast for the star trackers to maintain lock, while Override2 Mode is used to quickly converge the attitude solution once the star trackers are back in lock.

The attitude knowledge errors are fairly stable when the star trackers are enabled. The primary source of error is a slowly changing bias in the star tracker data. The star tracker exhibits small random errors from measurement to measurement that are an order lower than the bias error, classified as temporal error. Additional small errors can be incorporated occasionally when a star drops out of the field of view, requiring a change in the set of stars used for the attitude solution.

When slewing too fast to maintain star tracker lock, the ADCS system relies solely on gyroscopes for attitude determination. Since the gyroscopes measure angular velocity, a constant bias in the angular velocity measurement shows up as a drift error between the true and nominal attitudes. In addition, the gyroscopes exhibit scale factor and misalignment errors that grow at a rate proportional to the angular velocity of the spacecraft. These errors become an issue at closest approach, when the angular velocity peaks at 1 degree per second.




Figure 6. Typical Attitude Knowledge Error Profile

attitude knowledge error converges to a new star tracker bias level, sampled independently of the original bias. The values assumed for each of these attitude errors are discussed in the next section.

During the Hartley 2 flyby, the control system function of ADCS is configured to point at comet Hartley 2. ADCS uses AutoNav's OD solution to construct a control frame for the spacecraft attitude. Two body-fixed attitude frames were used at the Hartley 2 flyby: (1) Image Sun attitude and (2) Image Protect attitude. Both attitude frames orient the Z-axis (HRI boresight) along the spacecraft-to-comet vector and use a reference vector to orient the X and Y axes. For Image Sun, this vector is the spacecraft-to-sun vector, while for Image Protect the reference vector is

perpendicular to the flyby trajectory plane. Image Sun aligns the line axis of the camera with the Sun direction, while Image Protect aligns the pixel axis of the camera with the flyby trajectory plane.

VII. Monte Carlo Simulation

A flight-code-in-the-loop Monte Carlo simulation was used to assess AutoNav’s statistical performance at the Hartley 2 flyby. Although computationally more expensive, a Monte Carlo capability was preferred over a linear covariance analysis, since some error sources associated with the flyby do not have a Gaussian distribution. The observed center of brightness offset error relative to the comet center of mass, for example, is a product of several parameters, including instantaneous comet pole, phase, and shape relative to the sun direction as well as spacecraft position and attitude. While the position, pole, and phase have Gaussian distributions, the sun illumination manifests itself as a fixed bias and the attitude error exhibits different distributions depending on the current ADCS mode. A simulation capability is needed that samples errors from defined distributions in each case and applies them to form the appropriate dynamic and measurement models. A second reason for using a Monte Carlo approach is the ability to directly apply these models to the simulated image and attitude data that are processed by the flight code. Image processing and orbit determination are the two AutoNav flight software modules included and tested in the Monte Carlo simulation.

The error models sampled in the simulation include spacecraft-comet relative position and velocity errors, comet orientation, comet phase, and attitude knowledge errors. Each of these errors represents the residual between the “nominal states” and “true states”. The nominal states in the simulation represent the best model or estimate of the flyby environment, while the true states represent the true dynamics of the flyby environment.

Figure 7 depicts the flow of information through the Monte Carlo simulation. Each case of the simulation is initialized with a nominal comet-relative spacecraft trajectory, indicated as `relative_traj_0` in the figure. This trajectory represents the initial conditions used to seed AutoNav based on the final ground-based OD solution produced prior to the flyby. The nominal trajectory is queried at the simulation start time to obtain the nominal position and velocity initial conditions, indicated as `nominal_state_0`. To obtain the true initial conditions, the uncertainty distribution of the ground OD at the simulation start time is sampled, as indicated in Table 3, and applied to the nominal initial conditions. The downtrack component of this distribution is typically the largest since it is oriented along the approach asymptote and has limited observability in the camera frame during early approach. The true and nominal initial conditions are then input to a time loop that simulates the spacecraft flight from six hours prior to closest approach to two hours after closest approach. Times within the simulation are represented relative to the nominal encounter closest approach time using the prefix E- or E+. From this point in the simulation, the true trajectory is propagated open loop, as shown on the left side of Figure 7, while the nominal trajectory is continually updated by AutoNav as it estimates the trajectory based on observations, as shown on the right side of Figure 7. At select times defined in the simulation, image processing and orbit determination events are triggered to exercise the AutoNav system.

The performance of the ADCS system is modeled in the Monte Carlo simulation using an attitude knowledge offset between the ADCS attitude estimate and the true spacecraft attitude. The ADCS attitude estimate is modeled in the simulation by constructing either the Image Sun or Image Protect control frame from the latest AutoNav solution. This attitude is defined as the nominal attitude. The true attitude is obtained by applying error models to the nominal attitude based on the predicted star tracker and gyroscope performance shown in Table 2. These error models represent uncertainties in the ADCS attitude

Table 3. Sampled Initial Condition Errors

Error Model	3- σ Uncertainty	Distribution
Position Crosstrack	20 km	Gaussian
Position Downtrack	300 km	Gaussian
Velocity Crosstrack	5 cm/s	Gaussian
Velocity Downtrack	5 cm/s	Gaussian
Comet Pole RA	360 deg	Gaussian
Comet Pole DEC	180 deg	Gaussian
Comet Phase	360 deg	Gaussian

Table 2. Sampled Attitude Knowledge Errors

Error Model	3- σ Uncertainty	Distribution
Startracker Bias	300 urad, each axis	Uniform
Star Changeout	100 urad, each axis	Gaussian
Startracker temporal noise	50 urad, each axis	Gaussian
Gyro Drift	500 urad/hr, each axis	Gaussian
Gyro Scale Factor	183 ppm, each axis	Gaussian
Gyro Misalignment	80 urad, each axis	Gaussian

estimation process. The attitude knowledge errors are sampled at the start of each Monte Carlo case and applied at each picture time to calculate the true attitude from the nominal attitude. The true attitude is used to model the true observations recorded by the camera, while the nominal attitude is used by AutoNav to interpret those observations.

The simulation uses an ellipsoid shape model to simulate MRI observations of the comet. The generation of the observations relies on the comet shape model, comet pole, comet phase, comet rotation rate, spacecraft-comet relative position, spacecraft attitude, sun direction, and camera model. In order to test a wide variety of lighting conditions, the initial conditions for the comet pole and phase are randomly sampled to define a body-fixed frame. As shown in Table 2, the comet pole and phase are sampled from all possible combinations since no a priori knowledge exists for these parameters. The ellipsoid shape model is affixed to the body-fixed frame and propagated through the encounter using the comet rotation rate. The light incident on the ellipsoid is projected onto the MRI CCD plane to generate an optical observation. As evident in Figure 7, the generation of the image represents an observation of the true environment and therefore is a product of true position, true comet frame, and true MRI attitude.

Each MRI image is input to the AutoNav Image Processing Module along with the nominal camera attitude to interpret the observations and accumulate a measurement data arc. At select times, the AutoNav Orbit Determination

Module is called to produce a least-squares estimate of the trajectory by fitting the existing arc of measurements. The new trajectory updates the nominal state model for the next set of observations.

The results of the Monte Carlo simulation are recorded as residuals between the observed comet location in the MRI and the predicted, or computed, comet location in the MRI based on the latest AutoNav trajectory estimate. The computed measurement is derived by calculating the expected comet location in the MRI based on the AutoNav position estimate and ADCS attitude estimate at the time of the observation. By subtracting the computed measurement from the observed measurement, the

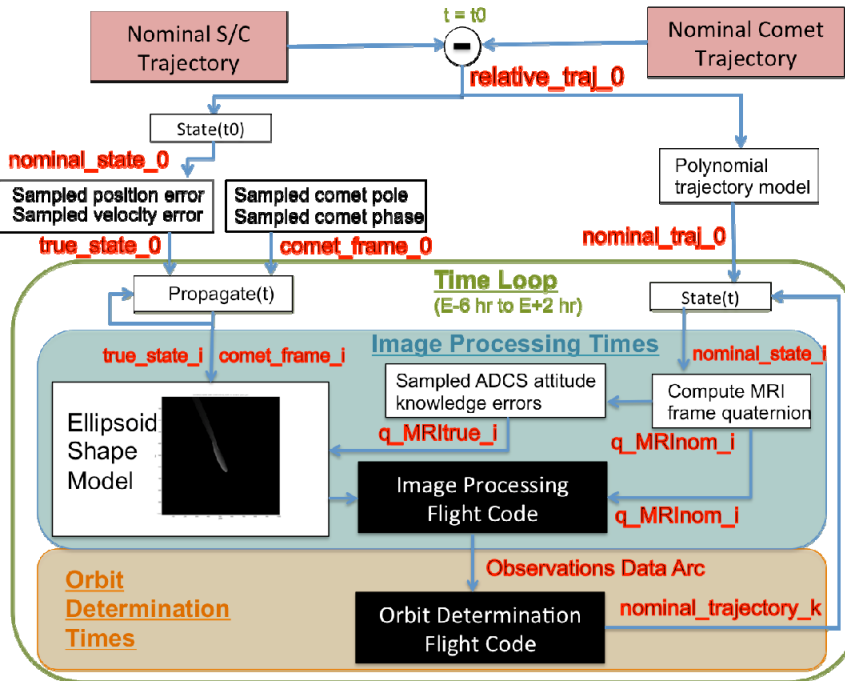


Figure 7. Monte Carlo Simulation Data Flow For Each Case

residuals provide insight into AutoNav's ability to track the comet in the MRI frame.

VIII. Operations Decisions

The results of the Monte Carlo simulation provided valuable information to the EPOXI project for encounter sequence decisions. Some decisions, such as filter uncertainties and data arclengths, could be configured by the AutoNav configuration files while other decisions, such as AutoNav event frequencies and flight system mode transitions, required system-level implementation in the encounter sequence. AutoNav simulation findings drove the following critical encounter planning decisions for approach, encounter, and departure.

1. Use AutoNav for the Hartley 2 flyby
2. Begin AutoNav sequence at E-50 minutes
3. Transition to IMAGE PROTECT attitude at E-35min
4. Set OD arclength to 8 minutes
5. Weight Centroid data at 100 pixels from E-7min to E+2min
6. Loosen the AutoNav filter a priori velocity sigma to 2.0 m/s

7. Transition to Override2 Mode at E+10min and Nominal Mode at E+18min
8. Revert to ground-based OD at E+50min

IX. Approach Phase

On approach to Hartley 2, a ground-based trajectory prediction was used by ADCS for spacecraft instrument pointing until AutoNav activation on the spacecraft. The first decision of encounter planning was to determine if AutoNav was needed for the flyby or if a ground-based trajectory would be sufficient. Figure 8 plots the observed growth of the 3- σ ground-based trajectory errors in camera pixels as the spacecraft nears closest approach. The error growth in pixel space obeys an inverse sine function that is inversely proportional to the range to the comet, peaking out at the time of closest approach when the minimum range of 700 km is achieved. The 3- σ uncertainty cloud peaks in the pixel axis at $\pm 60,000$ pixels, or 120 MRI fields of view, at closest approach, while the errors in the line axis grow to a maximum of $\pm 3,000$ pixels, or 6 MRI fields of view, at closest approach. The error projection onto the




Figure 8. Growth of Ephemeris Errors in MRI Frame vs. Time (3- σ)




Figure 9. Comet Size in MRI vs. Time

pixel axis is larger because both time-of-flight errors and in-plane crosstrack errors are projected in this direction while only out-of-plane errors are projected along the line axis. As observed in Figure 8, the comet nucleus would exit the HRI field at E-30 minutes and exit the MRI field at E-10 minutes, indicating the potential loss of all science from E-10 minutes to E+10 minutes without the use of AutoNav.

The second decision of encounter planning was determining when to initiate AutoNav usage on the spacecraft. The choice of a start time is ultimately a trade off of measurement signal strength versus the penalty of stale ephemeris errors. The measurement signal strength observed in the camera depends on many factors, including nucleus size, orientation, shape, and solar phase angle. At Tempel 1, AutoNav was configured to initiate comet imaging at E-2 hours^{4,5}. Although successful, the same configuration could not be assumed for Hartley 2 due to its smaller size, faster rotation rate, and elongated shape. With a solar phase angle near 90°, only half of the nucleus can be expected to have bright illumination, so the nucleus mean radius size in the MRI provides a measure of observability. Figure 9 shows the size of the Hartley 2 mean radius of 0.6 km in MRI pixels during the last two hours of approach. At E-2 hours, the illuminated portion of the nucleus would only be 0.7 pixel wide on average, less than the minimum 1 pixel size preferred for AutoNav. Since the comet size in MRI pixels increases as an inverse sine function with range to the comet, the signal strength could be dramatically improved by delaying the start of AutoNav imaging. However, delaying the AutoNav start time also increased the probability of science degradation due to a stale onboard ephemeris. A happy medium was found at E-50 minutes, when the pointing offset due to stale ephemeris errors was bounded at 80 pixels (3- σ) and the mean comet radius was projected to be 1.8 pixels in the MRI. The first AutoNav OD was set for E-42 minutes after a reasonable data arc of optical measurements was accumulated. The 80-pixel error bound safely kept the comet within the HRI field of view for IR spectroscopy and provided more than a factor of 5 margin against losing the comet in the MRI. During open loop

simulation runs, the comet was detected in the MRI at E-50 minutes in every case sampling random comet orientation and pole.

The third decision of approach was deciding when to transition from the Image Sun attitude frame to Image Protect attitude frame. Image Protect is the preferred attitude for a comet flyby since it orients the solar panels edge-on to the comet-relative velocity vector, reducing the risk of solar panel damage from a particle impact. It also aligns the pixel axis of the MRI with the trajectory plane, projecting the entire encounter slew onto the pixel axis. Since the Image Protect normal vector is defined by the cross product of the spacecraft-to-comet position and comet-relative velocity vectors, these vectors require a reasonable angular separation in order to provide a stable attitude frame definition. When the angular separation is small, the frame construction is highly sensitive to AutoNav OD updates, which update both the position and velocity vectors. Image Sun, on the other hand, is much less sensitive to AutoNav OD updates since the reference vector is the fixed sun direction. By the time of the first OD at E-42 minutes, the angular separation of the Image Protect reference vectors is approximately 1.3 degrees and increasing in the nominal trajectory.

At Tempel 1, the transition from Image Sun to Image Protect attitude control occurred at E-1 hour, long after AutoNav had begun providing OD updates^{4,5}. Since the first AutoNav OD was moved to E-42 minutes at Hartley 2, the decision to make was whether to keep the transition at E-1 hour or to move it after the start of AutoNav, where it could have an impact on AutoNav performance during a period of higher dynamics. The Monte Carlo simulation provided the capability to simulate the 3- σ statistics of normal vector updates to the Image Protect frame based on each AutoNav OD update. As shown in Figure 10, there is a potential for up to a 2-degree correction to the Image Protect normal vector at the first OD, adding extra stress to the ADCS system. However, by delaying the transition to Image Protect to E-35 minutes, the potential frame update was reduced to no greater than 0.04 degrees, resulting in a configuration with much greater stability. Therefore, based on the recommendations of AutoNav, the ADCS Image Sun to Image Protect transition was moved from E-1 hour to E-35 minutes in the encounter sequence.




Figure 10. 3- σ Normal Plane OD Updates

X. Encounter Phase

The encounter phase of the Hartley 2 flyby is characterized by a 180° slew through closest approach to track the comet. A variety of Monte Carlo cases were run to examine AutoNav performance under a variety of conditions, as shown in Table 4. These cases examined variations in comet shape model, attitude models, OD arc length, and filter weighting.

Figure 13 shows the results of the trade study, which drove decisions 4, 5, and 6, in residuals between the observed and predicted locations in MRI pixels. Case A represents a basic AutoNav configuration with a 20-minute OD arc length based on Tempel 1 experience⁴, no attitude knowledge errors applied, and a minimal center of brightness offset from the comet center of mass represented by a 200-meter radius round comet model.

Case B duplicates the Case A configuration, but incorporates the Hartley 2 ellipsoidal shape model. Examination of the differences in performance between Case A and Case B reveal the effects of a varying center of brightness offset on AutoNav performance. The residuals increase threefold in the pixel axis, but most of the increase in residuals occurs in the line axis, where most of the center of brightness offset is projected. While the line residuals in case A show a repeatable trend in each case that peaks at 5 pixels, case B exhibits a ± 80 pixel cloud at closest approach, representing how various orientations of the elongated comet can project different center of brightness offsets in the

Table 4. Encounter Monte Carlo Cases

Case	Comet Shape	Att Errors	OD Arc Length	Blobber Weight	Centroid Weight	Velocity Sigma
A	0.4km x 0.4km x 0.4km	Off	20 min	15 px	15 px	0.5 m/s
B	2.2km x 0.5km x 0.3km	Off	20 min	15 px	15 px	0.5 m/s
C	2.2km x 0.5km x 0.3km	On	20 min	15 px	15 px	0.5 m/s
D	2.2km x 0.5km x 0.3km	On	8 min	15 px	15 px	0.5 m/s
E	2.2km x 0.5km x 0.3km	On	8 min	15 px	100 px	0.5 m/s
F	2.2km x 0.5km x 0.3km	On	8 min	15 px	100 px	2.0 m/s

MRI. A second effect to note is the persistence of increased residuals in the line axis for several minutes after closest approach. This is a result of the AutoNav filter attempting to fit the shifting center of brightness signal to a dynamic trajectory model through closest approach. Since both pre-closest approach and post-closest approach observations are present in the OD data arc in the minutes after closest approach, AutoNav has difficulty locking on to the new observed center of brightness until post-closest approach observations outweigh pre-closest approach observations.

Case C applies attitude knowledge errors to the Case B configuration. The most pronounced effect is seen in the persistence of ± 40 pixel errors in both the pixel and line axes after closest approach, as observed in Figure 13. This

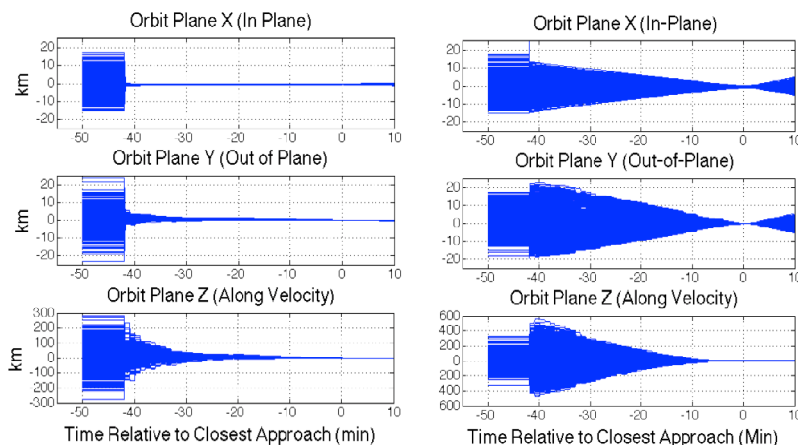


Figure 12. Comparison of Orbit Plane Position Errors without Attitude Errors (left) and with Attitude Knowledge Errors (right)




Figure 11. Orbit Plane Velocity Errors With Attitude Knowledge Errors

is partly an amplification of the combined pre-flyby and post-flyby data arc issues seen with the center of brightness offset. The effect in Case C is seen in both axes since the attitude knowledge errors are equally distributed in all axes while the center of brightness offset error is mostly contained in the line axis. The AutoNav filter initially absorbs the attitude knowledge bias errors as position corrections at the first OD update. Figure 12 shows how the AutoNav position estimate is adversely impacted by attitude knowledge errors, with 3- σ position errors at the first OD increasing from ± 1 km to ± 15 km crosstrack, from ± 4 km to ± 20 km normal, and from ± 200 km to ± 500 km downtrack.

Since a fixed position error grows exponentially in camera pixels with closing range to the comet, the filter must adjust its solution to match the constant bias profile observed in the camera from an attitude bias. In order for the attitude bias to match the observation profile of a position error, a velocity error must be incorporated that, together with the position error, appears as a linear constant offset in the camera frame. Figure 11 shows the 3- σ velocity error statistics in orbit plane coordinates. The normal axis absorbs the star tracker bias very quickly after the first OD, with a statistical spread of ± 500 cm/s. The in-plane axis initially absorbs only ± 60 cm/s of the star tracker bias, since downtrack cannot be discerned from crosstrack early in approach. Most of the growth in crosstrack velocity error occurs within E-10 minutes, when parallax between consecutive observations becomes much stronger.

Due to the issues mentioned above, a flexible and nimble filter configuration was desired that could offer quick adjustments to new observations. Case D examines the performance of Case C when the OD data arc is reduced from 20 minutes to 8 minutes. As seen in Figure 13, the residuals march down toward zero as the post-flyby observations increase and the pre-flyby observations decrease at each minute after closest approach. By E+8 minutes, the OD data arc is completely free of any pre-flyby data and exhibits maximum residuals of ± 10 pixels in the pixel axis and ± 20 pixels in the line axis. Based on this performance improvement, decision 4 was made to set the OD arc length to 8 minutes.

An effort was considered to speed the convergence of the residuals after closest approach to provide better pointing accuracy for the post-flyby IR scans. In order to rapidly converge the residuals after closest approach, a new weighting strategy was examined. The strategy involved significantly de-weighting the observation data around closest approach and reinstating tight weighting once the spacecraft was firmly on departure. The flight version of AutoNav provided the capability to independently adjust the nominal weights for Centroid and Blobber observations. In Case E, AutoNav was configured with a Blobber weight of 15 MRI pixels and a Centroid weight of 100 MRI pixels, applying more than a factor of 6 tighter weights on the Blobber data. In the encounter sequence,

the Blobber to Centroid transition was set at E-7 minutes, followed by a return back to Blobber at E+2 minutes, effectively de-weighting all observations from E-7 minutes to E+2 minutes. The motivation behind the choice of transition times was to have all tightly weighted pre-flyby data removed from the OD data arc by the time the first set of tightly weighted post-flyby data was incorporated. Given the order of magnitude difference in weighting, the expectation was that the OD at E+3 minutes would converge within a few pixels of the post-flyby observations. Despite having the desired effect on line residuals, this had the detrimental effect of increasing pixel residuals to over 300 at closest approach due to the de-weighting of the closest approach data, as seen in Figure 13. In this

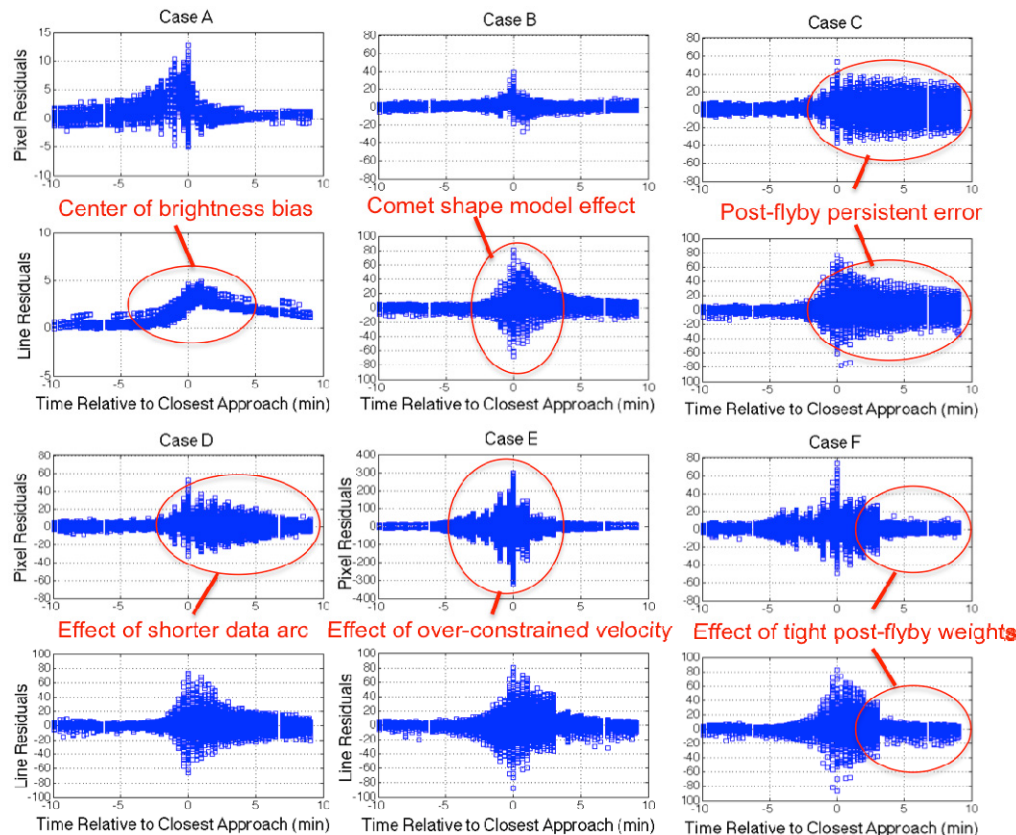


Figure 13. Observed-Predicted Residuals in MRI pixels

configuration, the comet could be at risk of falling outside of the MRI field of view.

The signatures seen in the case E pixel residuals around closest approach indicated an issue with the AutoNav velocity solution. As observed in Figure 13, the residuals snap down to zero at each OD, indicating good position determination, but the residuals grow to several hundred pixels between OD updates, indicating poor velocity determination. The root of the problem is due to the absorption of velocity updates in the AutoNav OD solution to match the attitude bias profile in the observations. While the line axis immediately absorbs the required velocity to match the attitude bias profile, the pixel axis cannot discern crosstrack errors from downtrack errors until parallax is observed between images within 10 minutes of closest approach. The AutoNav filter requires a data arc with multiple observations exhibiting noticeable parallax in order to distinguish crosstrack errors from downtrack errors. Initially, the attitude bias in the pixel axis is mostly interpreted as a significant error in downtrack position, since this direction has limited observability in the MRI and does not require a velocity change to match the profile in the camera. In other words, it is a simple location for the filter to alias the attitude bias error. As more parallax is observed between images, the filter begins to discern downtrack errors from crosstrack errors and the crosstrack velocity must be changed to match the attitude bias profile in the camera. In case E, the velocity sigma set in the AutoNav filter was 0.5 m/s, which was too constraining to allow a rapid crosstrack velocity change within 10 minutes of closest approach, especially with the de-weighted Centroid observations.

Case F loosened the AutoNav filter velocity sigma from 0.5 m/s to 2.0 m/s to allow more velocity mobility in the minutes around closest approach. As seen in Figure 13, the residuals in both axis snap down to within ± 5 pixels at the E+3 minute OD, with the exception of a few outliers. There is a small amount of degradation in the pixel axis,

where the residuals peak around 80 pixels at closest approach, but given the nucleus size of several hundred pixels at closest approach, this was less of an issue. Also, the de-weighting of observations around closest approach had the added benefit of reducing susceptibility to bright jet plumes or other observation disturbances around closest approach. Based on this improved performance, decisions 5 and 6 were made to de-weight the Centroid data from E-7min to E+2min and loosen the velocity sigma to 2.0 m/s.

XI. Departure Phase

The priorities on departure included effectively re-converging the attitude estimate after the closest approach slew and transitioning back to the ground-based trajectory. During the slew through closest approach, attitude knowledge errors accumulate due to the gyroscope scale factor, misalignment, and drift. By the end of the slew, the estimated attitude can be off from the true attitude by up to 1 milliradian. Most of these attitude errors are absorbed into the position or velocity solution by the AutoNav filter in order to quickly converge on the new center of brightness. When the star trackers reacquire lock after closest approach to stabilize and re-converge the attitude solution, the AutoNav solution exhibits a ringing effect as attitude knowledge errors are removed. This is observed

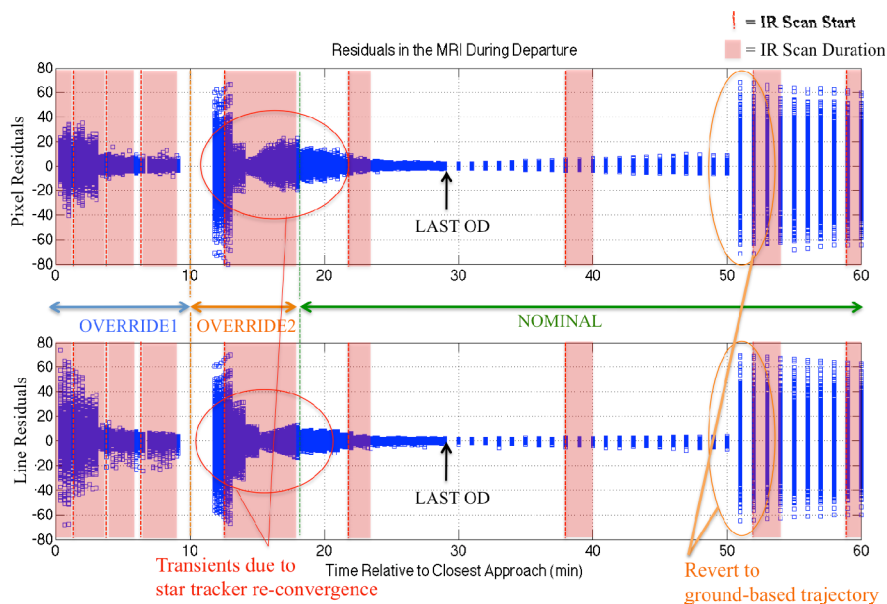


Figure 14. Observed-Predicted MRI Residuals During Departure

between E+12 minutes and E+26 minutes in Figure 14. As the attitude estimate converges, ADCS adjusts the spacecraft pointing based on its improved knowledge, exposing the attitude knowledge errors soaked up in the position estimate. The large residuals apparent by the exposed attitude knowledge errors cause AutoNav to correct its trajectory solution to match the observations. Due to the 8-minute OD arc length, the AutoNav solution is not fully stabilized until a full 8 minutes after attitude is converged. As the data arc marches forward accumulating more observations with the new converged attitude, the AutoNav solution removes the previous attitude knowledge errors from the solution and converges on a stable trajectory.

A major decision was determining when and how to converge the attitude solution to minimize the impact on AutoNav performance during the IR scans. Figure 14 plots the MRI residuals during departure with periods of IR scans of the comet nucleus highlighted in red. Decision 7 examined choices for incorporating star tracker data after closest approach to re-converge the attitude estimate. A transition from Override1 Mode to Override2 Mode was chosen to occur at E+10 minutes to quickly converge the attitude solution to the star tracker data. This time was chosen because AutoNav already had a scheduled 3-minute gap in observation data around E+10 minutes. The idea was to have a majority of the attitude transient occur while AutoNav was in the dark to minimize difficulties fitting the observation arc. Indeed, the attitude transient during the observation outage produced residuals of ± 60 pixels in the first AutoNav observations afterward. By the E+14 minute OD, the residuals were no greater than ± 20 pixels, meeting expectations for the long IR scan from E+12 minutes to E+18 minutes. E+18 minutes was chosen as the transition time from Override2 Mode to Nominal Mode to allow a full OD arc to be retired before settling on the

final attitude. Override2 Mode was used instead of a direct transition from Override1 to Nominal Mode in order to guarantee a stable attitude solution by the time of the final AutoNav OD at E+30 minutes.

By 8 minutes after the transition to Nominal Mode at E+26min, the OD data arc contained a full arc of observations with the final converged attitude. At this point, the OD solution was considered stable and the OD update process could be discontinued since observations become weaker as range to the comet increases. The final OD was set for E+30 minutes and that solution was used until the ground-based trajectory was reestablished at E+50 minutes. Reverting to the ground-based trajectory was set for E+50 minutes to maintain symmetry with approach and since the cloud of uncertainty was only reduced an additional ± 10 pixels by delaying to E+120 minutes. Reverting to the ground-based trajectory at E+50 minutes also provided once last hedge against complete AutoNav failure during the encounter.

XII. Flight Performance

Figure 15 shows the residuals from flight overlaid on the final predicted performance from the Monte Carlo Simulation. The performance in flight agreed closely with the results of the simulation. The first AutoNav OD makes a correction of approximately 12 pixels in the line axis and only 1 pixel in the pixel axis. A small drift of




Figure 15. AutoNav Residuals From Flight (Red) Overlaid on Monte Carlo Simulation Results

approximately 5 pixels in the pixel axis occurred 2 minutes before closest approach due to the de-weighting of the Centroid data. The maximum residual in flight was recorded about 1 minute after closest approach at 20 pixels in both the pixel and line directions. As predicted, this occurred due to the combination of pre-flyby and post-flyby measurements in the data arc. As expected, these residuals snap back into place once the tight Blobber weighting is applied a few minutes after closest approach. The flight residuals also show up to a 10-pixel transient as ADCS transitions through Override2 to converge the attitude solution. Overall, the flight performance fit nicely within the $3\text{-}\sigma$ statistical envelope simulated by the Monte Carlo simulation.

Figure 16 shows the position and velocity errors between the AutoNav-determined trajectory and a trajectory reconstruction of the flyby. The errors are plotted in the orbit frame, with the Z-axis aligned downtrack, the X-axis aligned in-plane crosstrack, and the Y-axis aligned out of plane. The top plot shows the full magnitude of the position errors, while the middle plot zooms in to examine the X and Y errors, and the bottom plot shows velocity errors. The initial downtrack position error from the ground-based trajectory was approximately 12.9 km, or 1.05 seconds time of flight, while the crosstrack error was 1 km and the normal error was 2.5 km. At E-42 min, the first OD update occurs, inflating downtrack error to approximately 80 km, while the crosstrack error moves to 4 km and the normal error adjusts to 1 km. While the normal errors walk linearly toward zero at closest approach, the crosstrack and downtrack errors follow a more indirect path as the filter attempts to differentiate the two

components. The Y-axis velocity error maintains a bias of 40-50 cm/s until closest approach, while the X-axis velocity error is initially near zero, but begins rapidly increasing as the filter distinguishes the downtrack error from crosstrack error around E-10 min. A signature is seen in the X and Z position errors around the Image Protect attitude transition time, as a data arc containing observations at two attitudes provides the filter some information to distinguish downtrack error from crosstrack error. By the E-6 min OD update, the all components of position error fall within the ± 3.5 km requirement, meaning the comet would still be tracked through closest approach in the MRI if AutoNav was disabled at the this time.




Figure 16. AutoNav Position and Velocity Errors in Orbit Frame

convergence to the post-flyby observations. At E+12 min, the position and velocity absorb the new star tracker bias errors and settle to final values by E+26 minutes. The Y-axis errors settle near zero, while the X and Z-axis errors settle to nonzero values, indicating errors in differentiating downtrack error from crosstrack error. All of these errors remain well within the envelope from the Monte Carlo simulation.

XIII. Conclusion

The EPOXI flyby of Hartley 2 exercised the use of AutoNav to capture high-resolution images of the comet nucleus at closest approach. In order to configure AutoNav for performance in a new dynamic environment, a Monte Carlo simulation was deployed. The simulation effectively characterized the performance of AutoNav at Hartley 2 and provided important insights about the effect of encounter design decisions on AutoNav performance. The AutoNav filter was configured to allow parameter flexibility as the spacecraft slewed through closest approach and quick adjustments after closest approach to correct the spacecraft pointing. AutoNav performance at the Hartley 2 flyby was a success, with AutoNav predicting the trajectory within 20 pixels in the MRI throughout the encounter. Images from Hartley 2 depicted a bi-lobed, peanut-shaped comet nucleus with several jets of CO₂ gas spewing ice particles from the nucleus, as shown in Figure 17.




Figure 17. Hartley 2 Images Before, During, and After Closest Approach

XIV. Acknowledgments

This work was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration. The authors would like to thank Steve Collins and Dustin Putnam for their contributions to understanding the ADCS system interactions with AutoNav. The authors would also like to thank the entire EPOXI flight team for their collaboration on an exciting and challenging mission.

XV. References

¹Riedel, J.E., Bhaskaran, S., Synnott, S.P., Desai, S.D., Bollman, W.E., Dumont, P.J., Halsell, C.A., Han, D., Kennedy, B.M., Null, G.W., Owen Jr., W.M., Werner, R.A., and Williams, B.G. "Navigation for the New Millenium: Autonomous Navigation for Deep Space-1," *Proceedings of the 12th International Symposium on Flight Dynamics*, Darmstadt, Germany, Jun. 1997.

²Bhaskaran, S., Riedel, J.E., Kennedy, B., Wang, T.C., "Navigation of the Deep Space 1 Spacecraft at Borrelly," AIAA paper 2002-4815, *AIAA/AAS Astrodynamics Specialist Conference and Exhibit*, Monterey, CA, August 5-8, 2002.

³Bhaskaran, S., Riedel, J.E., Synnott, S., "Autonomous Nucleus Tracking for the Comet/Asteroid Encounter: The Stardust Example", AAS 97-628; *AAS/AIAA Astrodynamics Specialist Conference*, SunValley, Idaho, 4-7 August 1997.

⁴Mastrodemos, N., Kubitschek, D.G., Synnott, S., "Autonomous Navigation for the Deep Impact Mission Encounter with Comet Tempel 1", *Space Science Reviews*, Vol. 117, 2005, pp. 95-121.

⁵Mastrodemos, N., Kubitschek, D.G., Werner, R.A., Kennedy, B.M., Synnott, S., Null, G.W., Riedel, J.E., Bhaskaran, S., Vaughan, A.T., "Autonomous Navigation for Deep Impact", AAS 06-177, *16th AAS/AIAA Space Flight Mechanics Conference*, San Diego, CA, January 22-26, 2006.